arXiv:1012.0560v2 [math.QA] 30 Aug 2011

TURAEV-VIRO INVARIANTS AS AN EXTENDED TQFT III

BENJAMIN BALSAM

Abstract. In the third paper in this series, we examine the Reshetikhin-Turaev and Turaev-Viro TQFTs at the level of surfaces. In particular, we show that for a closed surface Σ, $Z_{TV,\mathcal{C}}(\Sigma) \cong Z_{RT,Z(\mathcal{C})}(\Sigma)$, thus extending the equality of 3-manifold invariants proved in [Bal] to an equivalence of TQFTs. We also describe how to compute Turaev-Viro state sums on 3-manifolds with embedded ribbon graphs.

Introduction

In this paper we continue the work from [BK], [Bal] in which we generalized the Turaev-Viro state-sum invariant to manifolds with corners. This gave an extended Topological Quantum Field Theory (TQFT). Using this extended theory, we showed that for a closed 3-manifold $\mathcal{M}$, $Z_{TV,\mathcal{C}}(\mathcal{M}) = Z_{RT,Z(\mathcal{C})}(\mathcal{M})$, where $\mathcal{C}$ is a spherical fusion category, $Z(\mathcal{C})$ is its Drinfeld Center (which is modular) and $Z_{RT,Z(\mathcal{C})}$ is the Reshetikhin-Turaev invariant based on $Z(\mathcal{C})$. In this paper, we show that the TQFTs are isomorphic at the level of surfaces. Namely, if Σ is a closed surface, we show that there is a nautural isomorphism $Z_{TV,\mathcal{C}}(\Sigma) \cong Z_{RT,Z(\mathcal{C})}(\Sigma)$ of vector spaces. We also note that we actually get an equivalence of extended 3-2-1 theories if we impose mild restrictions on the allowed types of manifolds with corners.

It is easy to compute the dimensions of the above spaces:

$$\text{Dim}\, Z_{TV}(\Sigma_g) = \text{Dim}\, Z_{RT}(\Sigma_g) = \mathcal{D}^{2g-2} \sum_{i \in \text{Irr}(\mathcal{C})} d_i^{2-2g} \tag{0.1}$$

where $\mathcal{D}$ is the dimension of $\mathcal{C}$ and d_i is the dimension of simple object X_i. The vector spaces are therefore isomorphic, but this is not enough. We need to exhibit a natural isomorphism between the spaces.

The same question occurs in general when attempting to define any 2D modular functor. For example, in RT theory, one decomposes the surface Σ into a union of punctured spheres[1], evaluates Z_{RT} for each of them, and uses the gluing axiom to obtain $Z_{RT}(\Sigma)$. A priori, this appears to depend on the choice of decomposition of Σ. Refining earlier work by Moore and Seiberg, Bakalov and Kirillov [BK2000] proposed a set of *moves* (The "Lego-Teichmüller Game"). relating any two such decompositions. One can show that each of these moves corresponds to a certain natural isomorphism of vector spaces and any two "paths" between two chosen decompositions yield the same map. The space $Z(\Sigma)$ is therefore well defined.

In this paper we apply the results described above to TV theory. In [BK], we constructed an isomorphism

$$Z_{TV}(\Sigma) \cong \text{Hom}_{Z(\mathcal{C})}(\mathbf{1}, Y_1 \otimes \cdots \otimes Y_n) \tag{0.2}$$

This work was partially suported by NSF grant DMS-0700589 .

[1]Following [BK2000], we call this a *cut sytem*.

where Σ is an n-punctured sphere with boundary components labeled by $Y_1, \dots Y_n \in Irr(Z(\mathcal{C}))$. Notice that the space on the right of this equation is by definition $Z_{RT,Z(\mathcal{C})}(\Sigma; Y_1, \dots Y_n)$.

It is important to note that RT is defined using "pairs-of-pants" decompositions of surfaces, while TV is defined via cell decompositions. Since the latter is a local construction and the former is inherently nonlocal, comparing the two requires a natural way of passing between them. The solution is simple and is provided immediately by the surface parametrizations defined in [BK]. Using these, we can compute maps between TV state spaces that correspond to each of the moves between cut systems and check that such maps are compatible with the projector $H_{TV}(\Sigma) \longrightarrow Z_{TV}(\Sigma)$. Thus, we get a natural identification $Z_{RT}(\Sigma) \cong Z_{TV}(\Sigma)$.

In both RT and TV theory, once we know the value of the TQFT on a punctured sphere, we can use the gluing axiom to define $Z(\Sigma)$ for any surface. Thus, we get a well-defined vector space, up to natural isomorphism that depends only on the topology of Σ. We do this in each case by defining "intermediate" vector spaces which do depend on some choices[2] and demonstrating that we can identify all such spaces naturally. The key result in this paper is that we can pass between the theories in a natural way, so that $Z_{TV,\mathcal{C}}(\Sigma) \cong Z_{RT,Z(\mathcal{C})}(\Sigma)$ independent of any choices.

The paper is organized as follows. First, we briefly review the theory of parametrized surfaces from [BK2000]. Next, we examine the effect of passing between parametrizations on the associated TV state spaces. In particular, we show that each of the moves yields a natural map between state spaces, which under projection gives the same identification between vector spaces as that in RT. This establishes an equivalence of theories at the level of surfaces. Finally, we consider extended 3-manifolds with boundary and show that both theories give the same answer. Along the way, we explain how Turaev-Viro theory works for 3-manifolds with embedded ribbon graphs. The appendix contains some of the larger diagrams referenced in the paper.

This paper completes the program outlined in [BK] and continued in [Bal]. The reader is strongly encouraged to read these papers before this one, as they contain much prerequisite material.

Acknowledgments. The author would like to thank Sasha Kirillov for his help in writing this paper.

1. Surface Decompositions

In this section we briefly review the notion of a parametrized surfaces. For a complete exposition, see [BK2000]. Informally, a parametrization is a way of writing a surface Σ as the union on punctured spheres, together with a fixed identification of each punctured sphere with a *standard* sphere. The *standard sphere* with n punctures is defined formally as

$$S_{0,n} = \mathbb{CP}^1 \backslash \{D_1, \dots, D_n\}; D_j = \{z || z - z_j| < \epsilon\}, z_1 < \dots < z_n \tag{1.1}$$

where ϵ is sufficiently small so the boundary circles do not intersect. We also fix a point $p_i \in \partial D_i$. Note that we have fixed an ordering of the boundary circles, so we can refer to the set of boundary components by $\{\mathbf{1}, \dots, \mathbf{n}\}$.

[2]The parametrization in RT and the cell decomposition in TV.

Definition 1.1. An *extended surface* is a compacted oriented surface Σ , possibly with boundary, together with a fixed point p_α on each boundary component $(\partial\Sigma)_\alpha$.

Note that there are several other equivalent ways of defining an extended surface (See [BK2001]).

Definition 1.2. A *colored* extended surface is an extended surface together which a choice of label $Z_\alpha \in Z(\mathcal{C})$ for each marked point p_α.

We now give the main definition of this section. Let Σ be a colored extended surface.

Definition 1.3. A *parametrization* of Σ consists of

(1) A finite set C of non-intersecting simple closed curves on Σ such that $\Sigma\backslash C$ is of genus zero. We call C the set of *cuts* and fix a point on each cut.
(2) For each component Σ_a of $\Sigma\backslash C$, a homeomorphism $\psi : \Sigma_a \to S_{0,n_a}$

Two parametrizations are considered equivalent if they are isotopic [3]. There is a nice graphical way of describing parametrized surfaces. Namely, take the standard sphere with the graph as shown in Figure 1. This graph connects a single internal

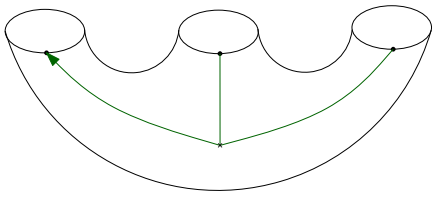


FIGURE 1. The graph on $S_{0,3}$

vertex to each of the points p_α fixed on the boundary and labels the edge connected to circle **1** by an arrow.

To depict a parametrization of any surface Σ, we draw the cuts on Σ. Then for each connected component Σ_α, we pull back the graph on the standard sphere by ψ_α to obtain a graph M_α on Σ_α. Clearly, such data are equivalent (up to isotopy) to specifying a parametrization and henceforth we will refer to a parametrization as a pair (C, M) where C is a set of cuts on Σ and $M = \cup_\alpha M_\alpha$. When possible, we will often draw the graphs M_α in the plane, ignoring the sufaces into which they are embedded. The reader should have no difficulty passing between such a graph and the surface it represents.

Definition 1.4. Let Σ be a parametrized sphere with n boundary components colored by $Z_1, \ldots, Z_n \in \mathrm{Irr}(Z(\mathcal{C}))$. We define the Reshetikhin-Turaev invariant of Σ to be

$$Z_{RT,Z(\mathcal{C})}(\Sigma; Z_1, \ldots, Z_n) = \mathrm{Hom}_{Z(\mathcal{C})}(\mathbf{1}, Z_1 \otimes \cdots \otimes Z_n) \tag{1.2}$$

More generally, we can defined the Reshetikhin-Turaev invariant for any colored, parametrized surface as follows. $\Sigma\backslash C$ is a union of genus zero surfaces with boundary, each equipped with parametrization inherited from Σ. Let $\Sigma_\alpha, \Sigma_\beta$ be two such components separated by a cut c. Then c corresponds to two boundary circles,

[3]Both the set of cuts and the homemorphisms of boundary components are considered up to isotopy

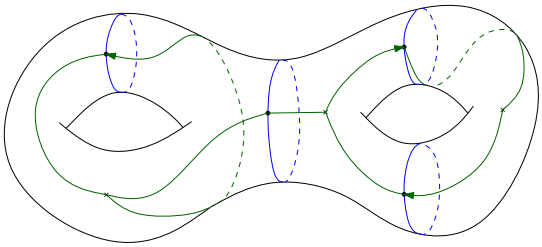


FIGURE 2. A parametrization of a genus two surface $\Sigma = S_{0,3} \sqcup S_{0,3} \sqcup S_{0,2}$. The blue lines are cuts and the green lines are graphs M_α

one on Σ_α and the other on Σ_β. We may color these components by assigning $Z \in \mathrm{Irr}(Z(\mathcal{C}))$ to one component and Z^* to the other.

Definition 1.5. Let (Σ, P) be a colored parametrized surface.

$$Z_{RT,Z(\mathcal{C})}(\Sigma, P) = \bigoplus_{Y_{1_\alpha},\dots Y_{n_\alpha}} \bigotimes_{\alpha} Z_{RT,Z(\mathcal{C})}(\Sigma_\alpha; Y_{1_\alpha}, \dots Y_{n_\alpha}) \tag{1.3}$$

where the product is over all connected components of $\Sigma \backslash C$, and we color all *newly created* boundary components as described above, summing over all possible colorings.

We will typically denote a simple object $Y_i \in Z(\mathcal{C})$ by its index i. In all that follows, i^* represents the dual object Y_i^*, which is also simple. To simplify formulas, many authors attempt to pick a function $f : Irr(\mathcal{C}) \to Irr(\mathcal{C})$ so that $Y_i^* = Y_{f(i)}$, but one should avoid doing this at all costs since it is often impossible to do so in a consistent manner (See [BK2001], Remark 2.4.2).

Example 1.6. Let Σ be the torus with one puncture and parametrization as shown on the left hand side of Figure 1 and boundary disk labeled by Y. Then $Z_{RT,Z(\mathcal{C})}(\Sigma) = \bigoplus_{i \in Irr(Z(\mathcal{C}))} Hom_{Z(\mathcal{C})}(\mathbf{1}, Y \otimes i \otimes i^*)$.

Now we describe a set of *moves* between parametrizations of a surface. As we'll see below, we can relate any two decompositions by a finite composition of these moves:

(1) The Z-move cyclically permutes the boundary components.
(2) The B-move *braids* one boundary component about an adjacent one.
(3) The F-move removes a cut. If a cut separates $S_{0,n}$ and $S_{0,m}$, deleting the cut gives a component homeomorphic to $S_{0,m+n-2}$ together with a graph inherited from the original components. Notice that we connect circle $\mathbf{1}$ from the one sphere to circle $\mathbf{m}$ of the other, thus resulting in a graph which inherits a natural ordering of boundary circles.
(4) The S-move interchanges meridians and longitudes of the punctured torus.

Theorem 1.7. *Let $A = (\Sigma, C, M)$ and $A' = (\Sigma, C', M')$ be two parametrizations of a surface. Then A and A' are related by a finite sequence of Z, B, F and S moves described above.*

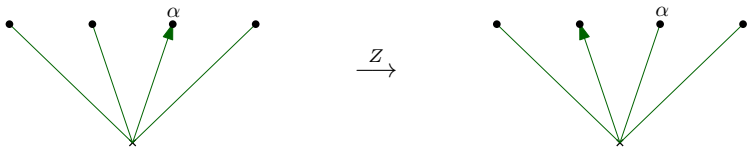


FIGURE 3. Z-move

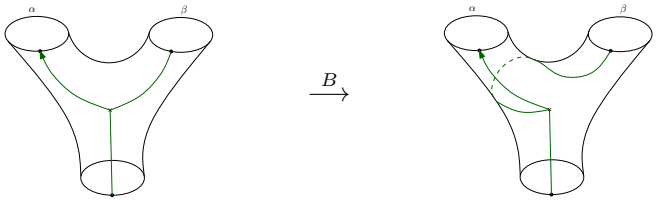


FIGURE 4. B-move

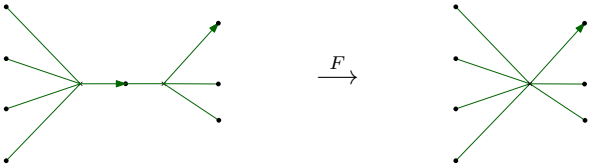


FIGURE 5. F-move

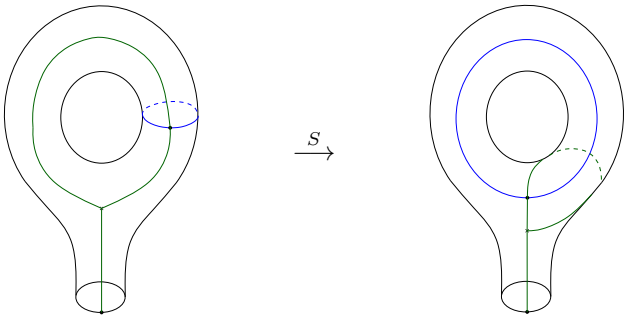


FIGURE 6. S-move

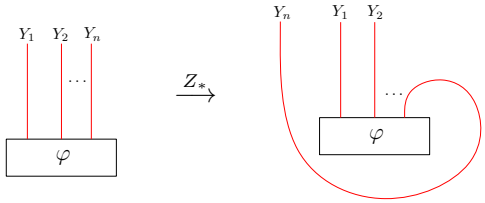
This result has its origins in conformal field theory. It was conjectured by Moore and Seiberg and rigorously proved in [BK2000]. It is a generalization of a result by Hatcher and Thurston ([HT1980]), which describes moves between surfaces decomposed into spheres, cylinders and pairs-of-pants, but doesn't take into account the full data of a parametrization. The theorem in [BK2000] does a lot more in fact: it provides a complete set of relations between the above moves, but we will not need this part explicitely.

These moves are important for defining a 2-dimensional (extended) modular functor $\mathcal{F}$. Given the vector space associated to the punctured sphere, one should be able to use the gluing axiom to describe $\mathcal{F}(\Sigma)$ for a surface of any genus. Different parametrizations should give naturally isomorphic vector spaces; one can check that this is so by verifying that it is true for each of the *simple* moves between

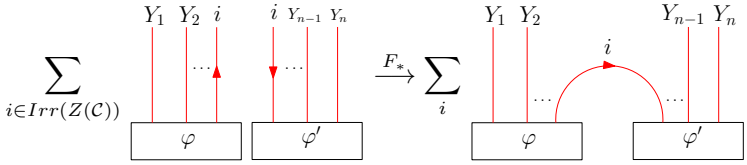
parametrizations. If P, P' are two parametrizations of a surface Σ related by a single Z, B, F or S move, we can explicitely describe the correspondence between associated vector spaces in RT theory:

Lemma 1.8. *Let P, P' be two parametrizations of a surface Σ and let $X : (\Sigma, P) \longrightarrow (\Sigma, P')$ be any composition of Z, B, F and S moves connecting P and P'. Then X induces an isomorphism $X_* : Z_{RT,Z(\mathcal{C})}(\Sigma, P) \xrightarrow{\cong} Z_{RT,Z(\mathcal{C})}(\Sigma, P')$. This isomorphism is independent of the choice of X. In terms of the generators,*

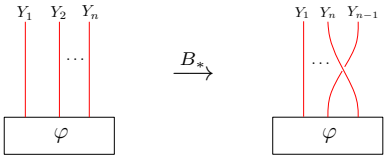
(1) *The Z-move corresponds to the rotation isomorphism:*
$\langle Y_1, \dots, Y_n \rangle \to \langle Y_n, Y_1, \dots, Y_{n-1} \rangle$



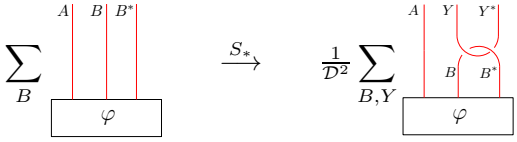
(2) *The F-move gives the composition isomorphism. That this is an isomorphism follows directly from semisimplicity.*



(3) *The B-move gives the braiding isomorphism*



(4) *The S-matrix gives multiplication by the S-matrix*



2. Parametrized surfaces and cell decompositions

In this section, we state and prove the main result of the paper: TV and RT theories assign the same vector space (up to natural isomorphism) surface Σ, which may have boundary. If $\partial\Sigma \neq \emptyset$, we fix a coloring of each boundary component $\partial\Sigma_i$ by $Z_i \in \mathrm{Irr}(Z(\mathcal{C}))$.

Given two cell decompositions Δ, Δ' of a surface Σ, there is a natural map

$$\Psi_{\Delta',\Delta} : H(\Sigma, \Delta) \longrightarrow H(\Sigma, \Delta') \tag{2.1}$$

obtained by computing a state sum on the cylinder $\Sigma \times I$, with a decomposition chosen to agree with Δ on $\Sigma \times 0$ and Δ' on $\Sigma \times 1$. Note, that this map does not depend on the choice of the internal decomposition. We will refer to this map as the *cylinder* map.

The cylinder map is not an isomorphism in general since the dimension of $H_{TV,\mathcal{C}}(\Sigma,\Delta)$ depends on the number of edges of Δ, but it is almost an isomorphism. More precisely, define the space $Z_{TV,\mathcal{C}}(\Sigma,\Delta) = Im(\Psi_{\Delta,\Delta})$. Then

$$\Psi_{\Delta,\Delta'} : Z_{TV,\mathcal{C}}(\Sigma,\Delta) \longrightarrow Z_{TV,\mathcal{C}}(\Sigma,\Delta') \tag{2.2}$$

is a natural isomorphism. We can refer to this space as $Z_{TV,\mathcal{C}}(\Sigma)$, since up to natural isomorphism it doesn't depend on the cell decomposition.

Given a parametrized surface Σ, there is a natural way to obtain a cell decomposition of Σ . We have a fixed collection of closed curves dividing Σ into the union of punctured spheres. These cuts become 1-cells in the cell decomposition. Further, for each punctured sphere thus obtained, we have a graph from our parametrization terminating at fixed points on the boundary circles. Each edge of this graph becomes a 1-cell and the points at which the 1-cells terminate become vertices. It is easy to see that these choices define a cell decomposition in the sense of [BK] [4]. We call the cell decomposition obtained in this way, the *associated* cell decomposition to parametrization P.

Recall that for a punctured sphere with standard cell decomposition (Figure 1), we have a projection $H_{TV,\mathcal{C}}(S^2) \xrightarrow{\pi} Z_{TV,\mathcal{C}}(S^2) \cong Z_{RT,Z(\mathcal{C})}(S^2)$. The associated inclusion map i can be described graphically: The normalization factors are chosen

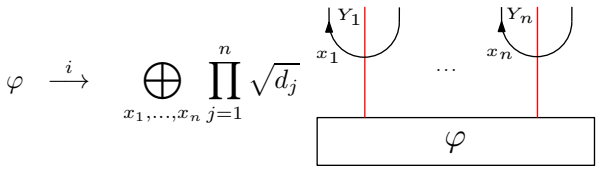


FIGURE 7. $i : Z_{RT,Z(\mathcal{C})}(S^2) \hookrightarrow H_{TV,\mathcal{C}}(S^2)$

to agree with that in [BK], so that $\pi \circ i = Id$.

We have two parallel notions in TV and RT theory. On the RT side, we have surface parametrizations and passing between any two parametrizations gives an isomorphism as described earlier in Lemma 1.8. On the TV, side, we have cell decompositions; passing between any two cell decompositions gives a natural isomorphism obtained from a cylinder as described earlier. The following theorem, which implies the main result in this paper, shows that these two notions are the same, up to projection.

Theorem 2.1. *Let P, P' be two parametrizations of a surface Σ with associated cell decompositions Δ, Δ' respectively. Then the diagram in Figure* 8 *commutes.*

Here, X_* is the map described in Lemma 1.8*, j is the map described in Figure* 7 *followed by projection to $Z_{TV,\mathcal{C}}(\Sigma)$and Ψ is the isomorphism described in* (2.2)*.*

Proof. To show the diagram commutes, we will verify that it does for each of the generators Z, B, F and S. The Z and B moves are essentially immediate, while the F and S moves require some work. Throughout the proof, our convention will be that diagrams of surfaces represent the vector spaces associated to them. In particular, a parametrized surface (Σ, P) represents $Z_{RT,Z(\mathcal{C})}(\Sigma, P)$ and a cell-decomposed surface (Σ,Δ) represents $Z_{TV,\mathcal{C}}(\Sigma,\Delta)$. In the diagrams below, we

[4]See, in particular, Figure 27

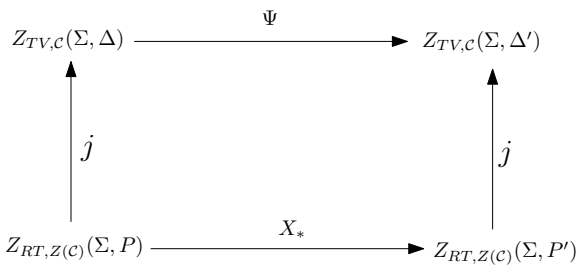


FIGURE 8.

have written Ψ from Figure 8 as the composition of several elementary steps for the reader's edification. We have moved several of the large diagrams to Section 4.

The Z-move. This follows directly from the natural isomorphism from Lemma 1.8(1).

The B-move. A proof of this fact may be found in [BK] (lemma 2.1), where we provide an explicit computation.

The F-move. We will show that the diagram in Figure 16 commutes . The arrow labeled F is the isomorphism described in Lemma 1.8, those labeled i are inclusion maps (Figure 7), and G is the gluing isomorphism at the level of state-spaces (Theorem 7.3, [BK]) . The other maps are all *cylinder* maps 2.1. Notice that this can be done in fewer steps, but the cylinder maps will be more difficult to realize. Figure 16 by contrast contains cylinder maps that are all easy to compute.

To check that this diagram commutes we begin with a vector in $Z_{RT,Z(\mathcal{C})}$ and proceed about the diagram in two ways. In Figure 17 we give the answer. The explicit computation at each stage left to the reader.

The S-move. We will show that the diagram in Figure 18 commutes. We have omitted some intermediate steps on the right side of the diagram as they are much the same as those on the left. Notice that the diagrams connected by the horizontal arrow labeled S are parametrized surface while the others are of cell-decomposed surfaces. We have chosen a convenient cell decomposition as the terminating point of the diagram which is easy to work with since there are simple maps α, β to this space which can be though of as contractions along edges u_1 and u_2 respectively (Figure 9). If we start on the bottom left of figure Figure 18 and proceed around

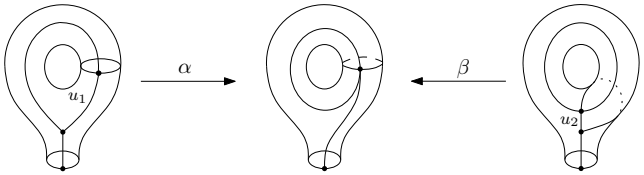
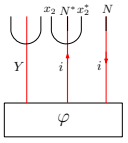


FIGURE 9. To identify the spaces on the left and the right, we use cylinder maps α, β to the space in the center and compare the images of these maps.

in two different ways, we get two vectors, φ_1, φ_2 in the same space as shown in Figure 10

$$\varphi_1 = \frac{1}{\mathcal{D}} \sum_{x_1,x_2,N} (d_{x_1} d_{x_2} d_N)^{\frac{1}{2}}$$ 

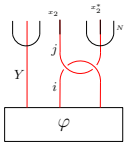
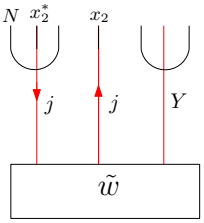
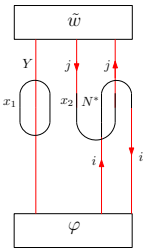
$$;\varphi_2 = \frac{1}{\mathcal{D}^3} \sum_{x_1,x_2,N} (d_{x_1} d_{x_2} d_N)^{\frac{1}{2}}$$ 

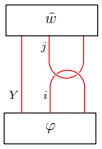
FIGURE 10.

We can easily verify that these vectors are the same by picking some vector w in the dual space and comparing the pairings $\langle \varphi_1, w \rangle$ and $\langle \varphi_1, w \rangle$. Let w be given by

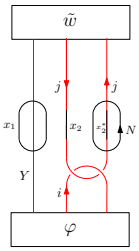
$$\frac{1}{\mathcal{D}} \sum_{i,N,x_1,x_2} (d_N d_{x_1} d_{x_2})^{\frac{1}{2}}$$ 

where $\tilde{w}$ is some vector in $\bigoplus_j Hom_{Z(\mathcal{C})}(\mathbf{1}, j^* \otimes j \otimes Y)$. Then

$$\langle w, \varphi_1 \rangle = \frac{1}{\mathcal{D}^2} \sum d_N d_{x_1} d_{x_2}$$ 

$$=$$ 

$$= \langle \tilde{w}, S_* \varphi \rangle$$

$$= \frac{1}{\mathcal{D}^4} \sum d_N d_{x_1} d_{x_2}$$ 

$$= \langle w, \varphi_2 \rangle$$ □

As an immediate consequence, we get

Theorem 2.2. *For any surface (possibly with boundary), we have a natural isomorphism*
$Z_{RT,Z(\mathcal{C})}(\Sigma) \cong Z_{TV,\mathcal{C}}(\Sigma)$.

3. Equivalence of Extended Theories

We conclude the paper by combining results in [BK], [Bal] and this paper to prove the following theorem.

Theorem 3.1. *Let $\mathcal{C}$ be a spherical fusion category. Then $Z_{TV,\mathcal{C}} \cong Z_{RT,Z(\mathcal{C})}$ as (3-2-1) TQFTs.*

We have already shown that the two TQFTs give the same answer for a closed 3-manifold with an embedded link [BK] (Theorem 4.8), and for a surface, possibly with boundary, as shown in Theorem 2.2. It remains to show that the theories give the same answer on any 3-manifold with corners.

Let us very briefly review the RT construction for 3-manifolds with corners. For more details, see [BK2001], [Tur1994]. Fix a spherical fusion category $\mathcal{C}$.

Definition 3.2. An extended 3-manifold $\mathcal{M}$ is an oriented PL 3-manifold with boundary, together with a finite collection of disjoint framed tubes $T_i \subset \mathcal{M}$.

An extended 3-manifold as described above is equivalent to a 3-manifold with an embedded framed tangle in the obvious way. We will use both descriptions interchangeably.

Notice that a tube T_i may terminate on $\partial\mathcal{M}$ in which case we call it an *open* tube, or it may close on itself, forming a solid torus, in which case we call it a *closed* tube.

Definition 3.3. A coloring of an extended 3-manifold $\mathcal{M}$ is a choice of color of simple object $Y \in Z(\mathcal{C})$ for each open tube T_i.

We wish now to generalize the famous theorem which states that any closed 3-manifold may be obtained from S^3 via surgery along a framed link.

Definition 3.4. A framed link with *coupons* is a framed link where components are allowed to coincide at multivalent vertices, called *coupons*. We often draw coupons as rectangles instead of vertices (Figure 11).

Given $L \subset \mathbb{R}^3$, an oriented, framed link with coupons, we can color L as follows. As before, assign to each edge, a simple object $Y \in Z(\mathcal{C})$. To each coupon C_i assign a morphism $\varphi_i \in W_i \equiv \mathrm{Hom}_{Z(\mathcal{C})}(\mathbf{1}, Z_1^{\epsilon_1} \otimes \cdots \otimes Z_n^{\epsilon_n})$, where $Z_1 \dots Z_n$ are the colors of the edges incident the coupon in clockwise cylic order, and $\epsilon_i = 1$ if the strand labeled by Z_i is oriented away from the coupon and -1 otherwise.

As shown in [Tur1994], we can evaluate such a link $L \subset \mathbb{R}^3$ to get a number $Z_{RT}(L) \in \mathbb{C}$ in a way that is invariant under isotopy of L. Further, since $\mathcal{C}$ (and hence $Z(\mathcal{C})$) is a spherical category, we can actually view L as lying in S^3.

Equivalently, if we leave the coupons of L uncolored, this construction gives a vector $v \in \bigotimes_i W_i^*$, where the tensor product is over all coupons in L and W_i is the Hom-space associated to coupon C_i. Thus, such a link with uncolored coupons gives a vector space V and a vector $v_L \in V$. Both can be seen to be invariant under isotopy of $L \in S^3$.

We are most interested in a particular type of oriented, framed link with coupons:

Definition 3.5. A *special* link X is a framed link with coupons such that some of the link components and coupons are colored by objects and morphisms, respectively, such that the following conditions hold:

- Any uncolored link component is either an annulus, or has both ends on the same uncolored coupon, in which case they are required to be adjacent to one another.
- Uncolored coupons are all of the form shown in Figure 11

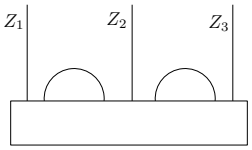


FIGURE 11. A uncolored coupon C in a special link X. Colored strands have a single end terminating on C. Uncolored strands have both ends terminating on C. Further, the ends are adjacent to one another.

Theorem 3.6. *Let $\mathcal{M}$ be an extended 3-manifold as in* [BK]. *Then $\mathcal{C}$ may be obtained from S^3 via surgery along some special link $X \subset S^3$, where we define surgery along X by*

$$M_X = M_L \backslash \bigcup_i T(C_i) \tag{3.1}$$

Here, we do ordinary surgery along all annular link components L, giving M_L, and remove handlebodies $T(C_i)$ which are tubular neighborhoods of uncolored coupons C_i, as shown in Figure 12.

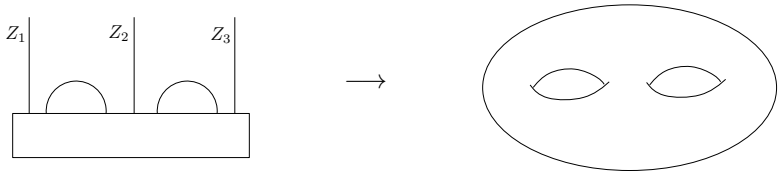


FIGURE 12. A coupon C determines a handlebody $T(C)$ by taking a tubular neighborhood of C and its uncolored strands. The genus of $T(C)$ equals the number of uncolored strands incident to C. The colored strands determine extra data attached to this handlebody, namely marked points and tangent vectors on $\partial(T(C))$ (not pictured), and are important in definining Reshetikhin-Turaev theory. For more details, see [BK2001].

Using the surgery description of an extended 3-manifold $\mathcal{M}$, we can easily define the RT invariant for such a manifold. Namely, we express $\mathcal{M}$ as the result of surgery along a *special* link $X \subset S^3$, and define

$$Z_{RT}(\mathcal{M}) \equiv Z_{RT}(X) \tag{3.2}$$

where $Z_{RT}(X)$ is obtained by evaluating the special link X, summing over all possible colorings of unlabeled link components, and using the convention that whenever we color an unlabeled component by simply object $Y \in Z(\mathcal{C})$, we multiply by d_Y, its categorical dimension. As noted above, if X has any uncolored coupons, then $Z_{RT}(X)$ is a vector, not a number.

We can also define Turaev-Viro theory on manifolds with embedded special links.

Definition 3.7. Let $\mathcal{M} = S^3_X$ be the 3-sphere with a special link X inside. Then

$$Z_{TV,\mathcal{C}}(\mathcal{M}) \equiv Z_{TV,\mathcal{C}}(\mathcal{M}') \tag{3.3}$$

where $\mathcal{M}'$ denotes the manifold with boundary obtained by removing tubular neighborhoods of each coupon.

Lemma 3.8. *Let $\mathcal{N}$ be a handlebody of genus g. Then $Z_{TV,\mathcal{C}}(\mathcal{N}) = Z_{RT,Z(\mathcal{C})}(\mathcal{N})$.*

This equality is to be interpreted as follows: Under the canonical isomorphism $Z_{TV,\mathcal{C})}(\Sigma) \cong Z_{RT,Z(\mathcal{C})}(\Sigma)$, where $\Sigma = \partial\mathcal{N}$, the two sides of the equation are identified.

Proof. As computed in [BK2001] (Example 4.5.3),

$$Z_{RT,Z(\mathcal{C})}(\mathcal{N}) = (id : \mathbf{1} \to (\mathbf{1} \otimes \mathbf{1})^{\otimes g}) \in \bigoplus \mathrm{Hom}(1, Z_1 \otimes Z_1^* \otimes \cdots \otimes Z_g \otimes Z_g^*). \tag{3.4}$$

. One shows that this equals $Z_{TV,\mathcal{C}}(\mathcal{N})$ by explicit state-sum computation. For $g = 1$, this can be deduced from [Bal], (Lemma 3.2). The general case is left to the reader as an exercise. □

Our goal is to show that $Z_{TV,\mathcal{C}}(\mathcal{M}) \cong Z_{RT,Z(\mathcal{C})}(\mathcal{M})$ for any extended 3-manifold $\mathcal{M}$. The idea is to convert our extended 3-manifold to a closed 3-manifold $\mathcal{N}$ (possibly with a link inside) by gluing handlebodies of appropriate genus to each component of $\partial\mathcal{M}$.

For simplicity, assume $\partial\mathcal{M}$ has a single component of genus g. Given a vector $\psi \in Z_{TV,\mathcal{C}}(\partial\overline{M})$, we can try to find a handlebody $\mathcal{H}_g$ with an embedded colored tangle, such that $Z_{TV,\mathcal{C}}(\mathcal{H}_g) = \psi$. By the gluing axiom, we get

$$\langle Z_{TV,\mathcal{C}}(\mathcal{M}), \psi\rangle = \langle Z_{RT,Z(\mathcal{C})}(\mathcal{M}), \psi\rangle$$

If we can do this for any ψ, then we are done.

Unfortunately, it is almost never possible to produce such a handlebody, even when $g = 0$. For example, if Σ is the 3-punctured sphere with boundary components labeled by Z_1, Z_2, Z_3, the space $Z_{TV,\mathcal{C}}(\Sigma)$ may be quite large, but there are no extended 3-manifolds with boundary Σ. Indeed Σ is cobordant to $\emptyset$ if and only if it has an even number of punctures.

One approach is to redefine $Z_{TV,\mathcal{C}}(\Sigma)$ as the vector space generated by $\{Z_{TV,\mathcal{C}}(\mathcal{M})\}$, where $\mathcal{M}$ ranges over all extended 3-manifolds with $\partial\mathcal{M} = \Sigma$. The theorem then follows from the above argument and some minor details, which we omit. This approach certainly works, but it is undesirable, e.g. it assigns a zero-dimensional vector space to any surface with an odd number of boundary circles.

A better approach is to use coupons.

Definition 3.9. Let $\mathcal{K}$ be the 3-ball with an special link inside as shown in Figure 13. We define

$$Z_{TV,\mathcal{C}}(\mathcal{K}) = \psi \in Hom_{Z(\mathcal{C})}(\mathbf{1}, Z_1 \otimes Z_2 \otimes Z_3 \otimes \cdots \otimes Z_N) \tag{3.5}$$

Notice that if we removed from $\mathcal{K}$ a tubular neighborhood of the coupon, we would be left with a cylinder over the N punctured sphere. We think of the coupon as a handlebody $\mathcal{H}$ which satisfies $Z_{TV,\mathcal{C}}(\mathcal{H}) = \psi$.

Also note that we have defined the value of $Z_{TV,\mathcal{C}}(\mathcal{K})$. It is not a result we can deduce from standard Turaev-Viro theory. However, it is consistent with the rest of our theory. In particular, we can treat $\mathcal{K}$ as an ordinary extended 3-manifold and use the gluing axiom, the graphical calculus descibed in [BK]. In particular, if we view the braiding isomorphism, the cup and the cap ([Bal], Section 2) as special examples of coupons, we get the same result as in Definition 3.9.

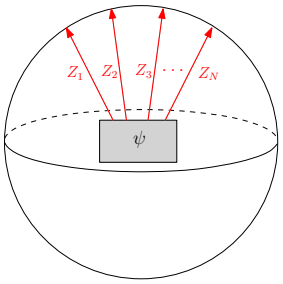


FIGURE 13. The manifold $\mathcal{K}$ is the 3-ball B^3 with a special link consisting of a single coupon and N strands connecting the coupon to ∂K. The coupon is labeled by some morphism $\psi \in \mathrm{Hom}_{Z(\mathcal{C})}(\mathbf{1}, Z_1 \otimes \cdots \otimes Z_N)$.

Theorem 3.10. *Let $\mathcal{M} = S^3_X$ be the 3-sphere with a special link X inside. Then*

$$Z_{TV,\mathcal{C}}(\mathcal{M}) = Z_{RT,Z(\mathcal{C})}(\mathcal{M}) \tag{3.6}$$

Proof. A special case of this theorem is proved in [Bal] (Theorem 2.3), where the result is proved if X is a colored link (with no coupons). If every coupon of X has the same form as that in Figure 13 (every strand that is incident to a coupon is colored and touches the coupon exactly once), the theorem follows immediately from Theorem 2.3 in [Bal], Definition 3.9 and the gluing axiom.

The situation is slightly more complicated if there are coupons of X that have uncolored strands (see Figure 11). We could try to come up with an analogous defintion to Definition 3.9 for the more complicated coupons, but in this case, the result follows from the state sum formula and Definition 3.9.

Let $\mathcal{H}$ be the extended 3-manifold shown in Figure 14. $\mathcal{H}$ is a cobordism, so by standard theory, it gives a linear map

$$Z_{TV,\mathcal{C}}(\mathcal{H}) : Z_{TV,\mathcal{C}}(S^2, Z_1, Z_1^*, Z_2, Z_2^*) \longrightarrow Z_{TV,\mathcal{C}}(\Sigma_2) \tag{3.7}$$

Notice that the left hand side of (3.7) is naturally a subspace of $Z_{TV,\mathcal{C}}(\Sigma_2)$.

Lemma 3.11. *The map defined above is given by*

$$Z_{TV,\mathcal{C}}(\mathcal{H}) = Id : Z_{TV,\mathcal{C}}(S^2, Z_1, Z_1^*, Z_2, Z_2^*) \longrightarrow Z_{TV,\mathcal{C}}(\Sigma_2) \tag{3.8}$$

where by Id, we mean the identification of the domain with its image under the identity map. If we sum up over all possible colorings of the strands, the two sides of the equation are naturally isomorphic, and we get the identity map.

Proof. We decompose $\mathcal{H}$ as shown in Figure 15.

The result follows immediately from ([BK], Example 9.2), where the computation is done in detail. □

An analogous result holds for a handlebody of any genus.

Now we can use Lemma 3.11 to finish proving Theorem 3.10. Suppose X contains a coupon C with uncolored strands beginning and ending on C (see Figure 11 for an example.) Let T(C) be a tubular neighborhood of C, as described earlier. It is a handlebody of some genus g. By definition,

$$Z_{TV,\mathcal{C}}(S^3_X) = Z_{TV,\mathcal{C}}(S^3_X \backslash T(C)).$$

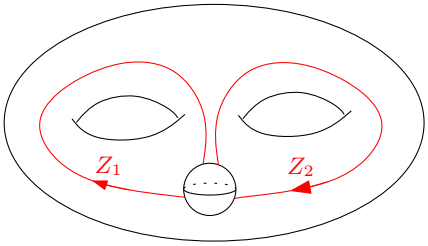


FIGURE 14. The extended manifold $\mathcal{H}$ is obtained by taking a handlebody of genus 2, removing a 3-ball and embedding a colored tangle inside as shown. Thus, $\mathcal{H}$ is a cobordism between the sphere with 4 holes and Σ_2, a suface of genus 2.

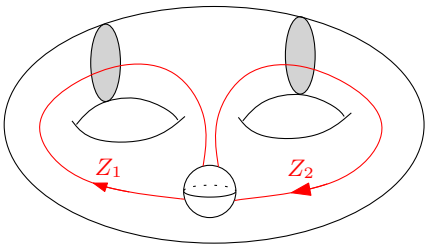


FIGURE 15. The decomposition of $\mathcal{H}$ is given by cutting along the grey disks. The complement of these disks is a cylinder over the sphere with 4 holes. We choose the same cell decomposition (not pictures) for this cylinder as in [BK], Example 9.2.

Using Lemma 3.11, we can glue an extended manifold $\mathcal{H}$ to $S^3_X \backslash T(C)$. (Here, we sum up over all colorings of strands inside $\mathcal{H}$.) Up to natural isomorphism $Z_{TV,\mathcal{C}}(\mathcal{H})$ is the identity map, so gluing it to $S^3_X \backslash T(C)$ does not change the value of $Z_{TV,\mathcal{C}}(S^3_X \backslash T(C))$.

Our new manifold may be described as $S^3_{X'}$, where X' is the same as X except the coupon C is replaced by a coupon C with no uncolored strands incident to it (so $T(C')$ has genus zero). Repeating this, we reduce X to a special link all of whose coupons have no uncolored strands incident to them. But we already know the theorem to be true in this case! □

We know from before that any extended 3-manifold may be obtained from S^3_X by doing surgery along the annular components of X, and removing tubular neighborhoods of coupons of X. Combining Theorem 3.10 with the surgery formula from [Bal] (Lemma 4.7) gives a proof of Theorem 3.1.

4. DIAGRAMS

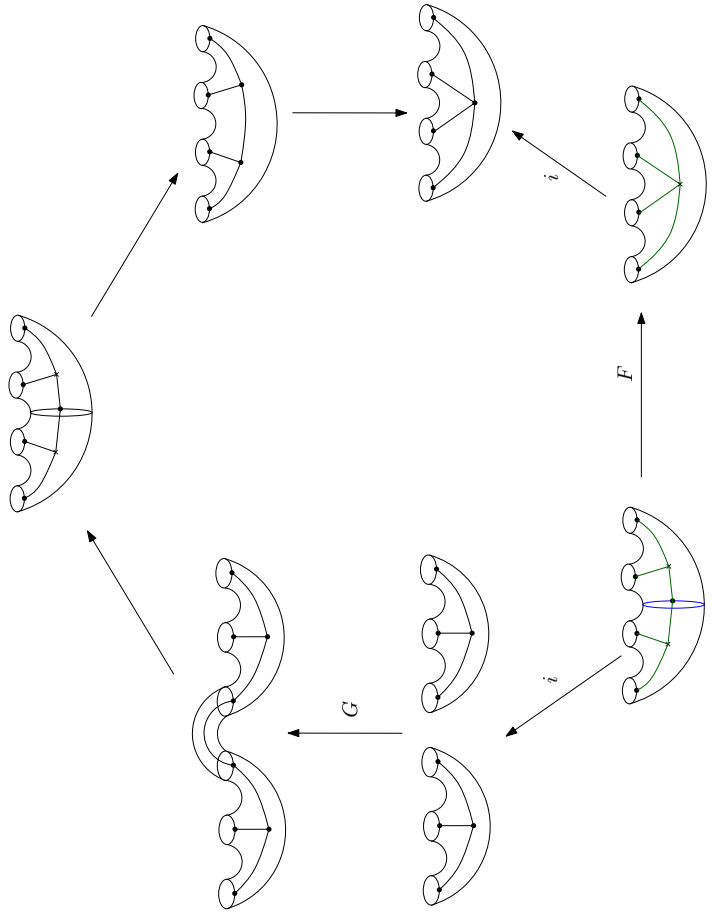


FIGURE 16. The F-move fuses together two spheres along a boundary component. In terms of parametrized surfaces this is realized by simply removing a cut separating the spheres. At the level of cell decompositions we want to identify the result with the standard sphere decomposition (Figure 1). We include several intermediate steps to make the computation more transparent.

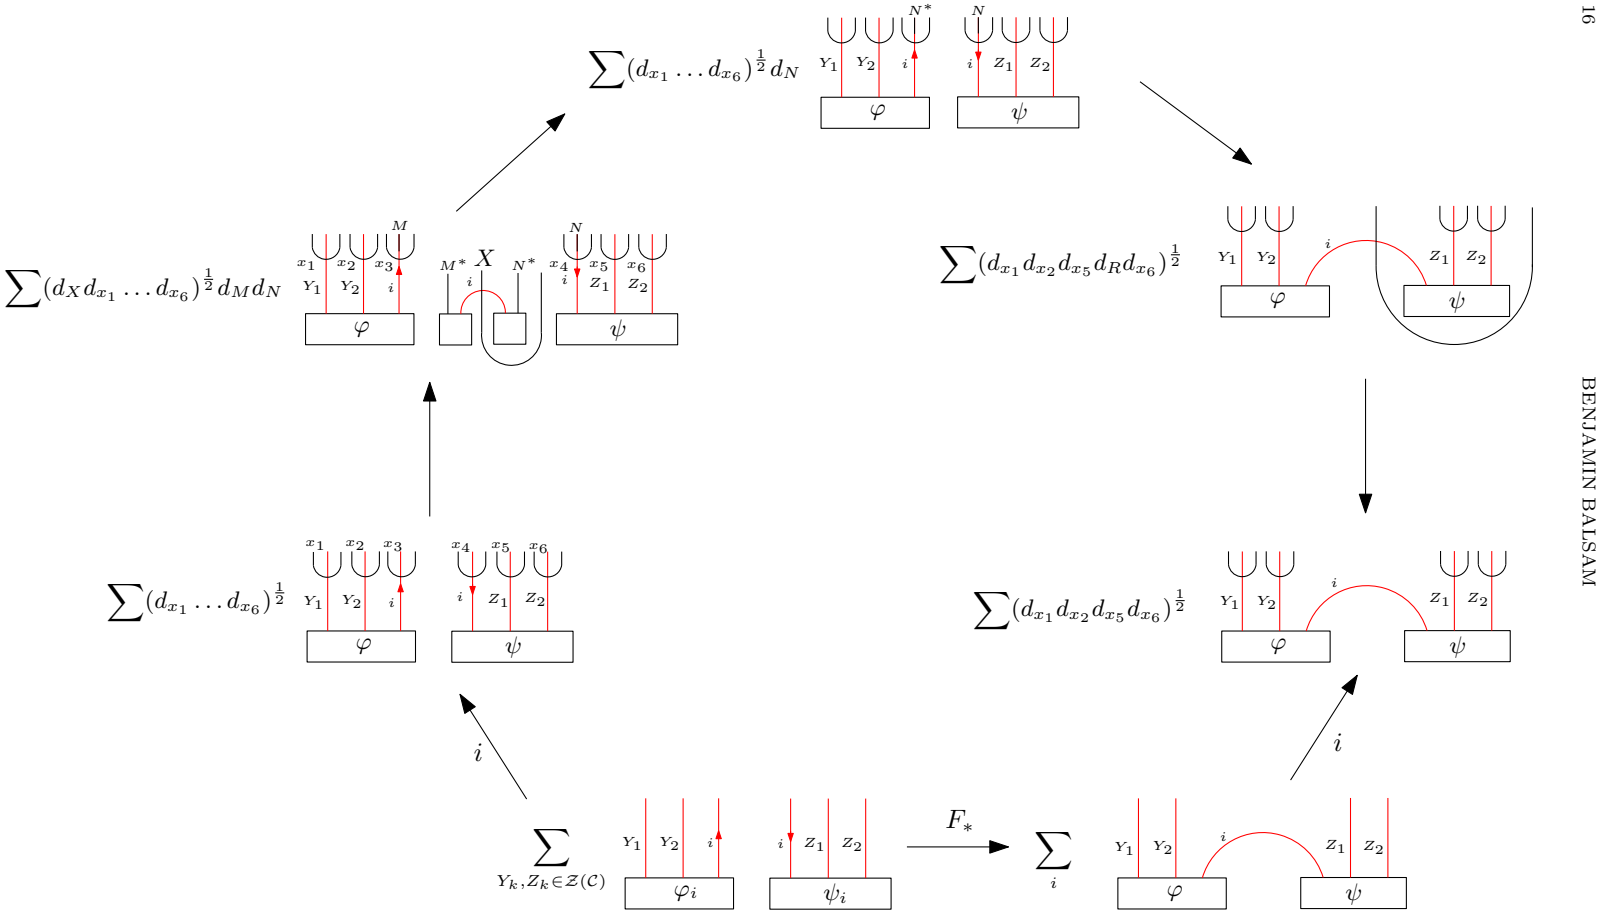


FIGURE 17. A verification that Figure 16 commutes. All unlabeled arrow represent cylinder maps. We start on the lower left and proceed in two ways around the diagram.

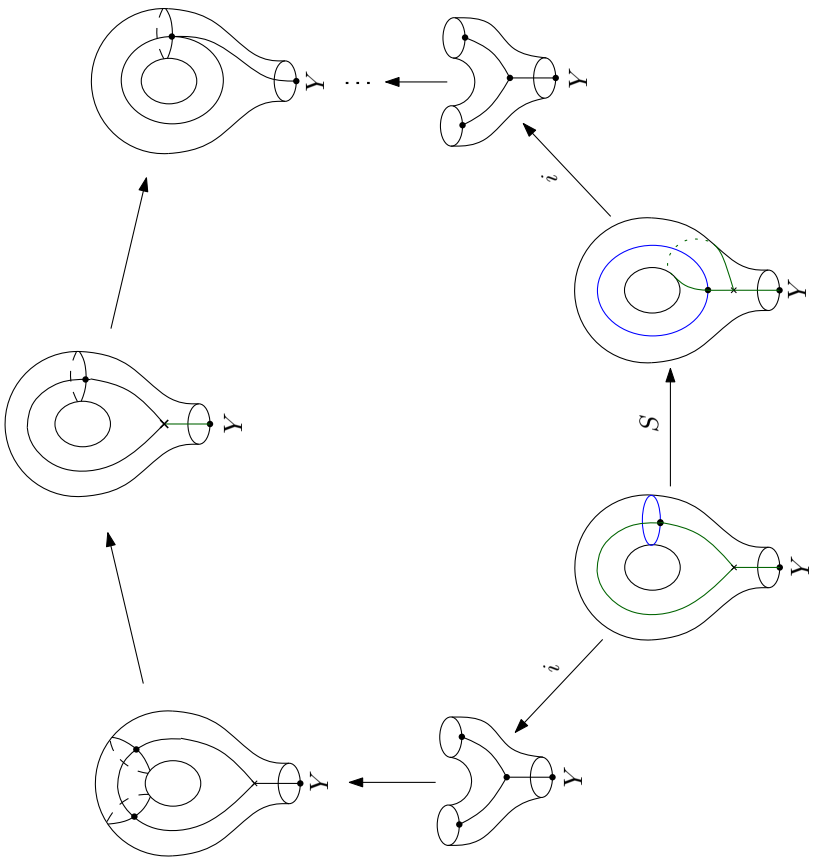


FIGURE 18. The S-move interchanges meridians and longitudes of the 1-punctured torus. Thus, the cut (blue) and the parametrizing graph (green) exchange places under the application of S.

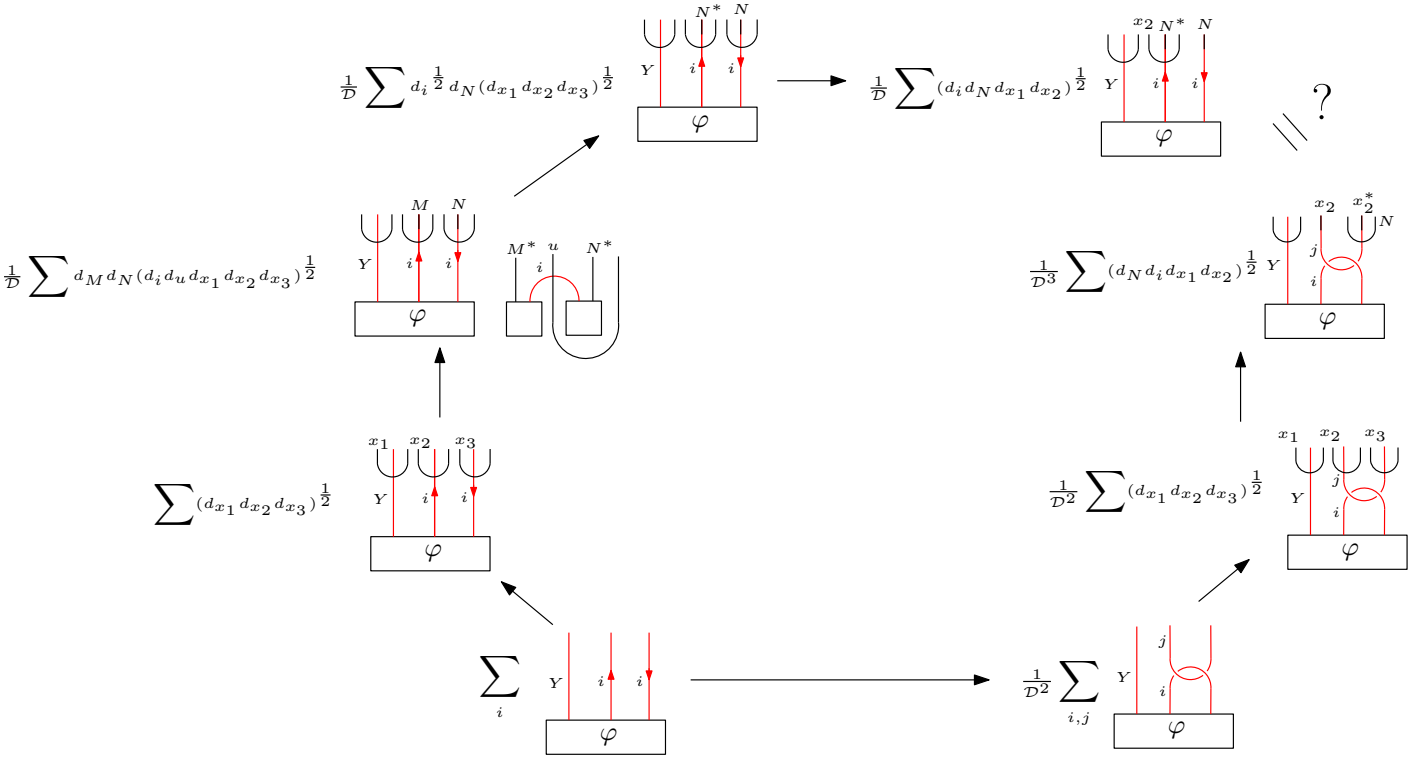


FIGURE 19. A demonstration that Figure 18 commutes. Notice that when we proceed around the diagram we get two different pictures (separated by " $\stackrel{?}{=}$ "), but can verify that they represent the same vector.

References

[BK] Benjamin Balsam and Alexander Jr. Kirillov, *Turaev-Viro Invariants as an Extended TQFT*, available at `arXiv:1004.1533`.

[Bal] Benjamin Balsam, *Turaev-Viro Invariants as an Extended TQFT II*, available at `arxiv:1004.1533`.

[BK2001] Bojko Bakalov and Alexander Kirillov Jr., *Lectures on tensor categories and modular functors*, University Lecture Series, vol. 21, American Mathematical Society, Providence, RI, 2001. MR1797619 (2002d:18003)

[BK2000] ———, *On the Lego-Teichmüller game*, Transform. Groups **5** (2000), no. 3, 207–244.

[BW1996] John W. Barrett and Bruce W. Westbury, *Invariants of piecewise-linear 3-manifolds*, Trans. Amer. Math. Soc. **348** (1996), no. 10, 3997–4022, DOI 10.1090/S0002-9947-96-01660-1. MR1357878 (97f:57017)

[ENO2005] Pavel Etingof, Dmitri Nikshych, and Viktor Ostrik, *On fusion categories*, Ann. of Math. (2) **162** (2005), no. 2, 581–642, DOI 10.4007/annals.2005.162.581. MR2183279 (2006m:16051)

[HT1980] Allen Hatcher and William Thurston, *A presentation for the mapping class group of a closed orientable surface*, Topology **19** (1980), 221-237.

[Lur] Jacob Lurie, *On the classification of topological quantum field theories*, available at `http://www-math.mit.edu/~lurie/`.

[Müg2003a] Michael Müger, *From subfactors to categories and topology. I. Frobenius algebras in and Morita equivalence of tensor categories*, J. Pure Appl. Algebra **180** (2003), no. 1-2, 81–157, DOI 10.1016/S0022-4049(02)00247-5. MR1966524 (2004f:18013)

[Müg2003b] ———, *From subfactors to categories and topology. II. The quantum double of tensor categories and subfactors*, J. Pure Appl. Algebra **180** (2003), no. 1-2, 159–219, DOI 10.1016/S0022-4049(02)00248-7. MR1966525 (2004f:18014)

[Tur1994] V. G. Turaev, *Quantum invariants of knots and 3-manifolds*, de Gruyter Studies in Mathematics, vol. 18, Walter de Gruyter & Co., Berlin, 1994. MR1292673 (95k:57014)

[TV] V. G. Turaev and Alexis Virelizier, *On two approaches to 3-dimensional TQFTs*, available at `arXiv:1006.3501`.

[TV1992] V. G. Turaev and O. Ya. Viro, *State sum invariants of 3-manifolds and quantum $6j$-symbols*, Topology **31** (1992), no. 4, 865–902, DOI 10.1016/0040-9383(92)90015-A. MR1191386 (94d:57044)

Department of Mathematics, SUNY at Stony Brook, Stony Brook, NY 11794, USA
E-mail address: `balsam@math.sunysb.edu`
URL: `http://www.math.sunysb.edu/~balsam/`